

MESSAGE RING IN A SWITCHING NETWORK

Introduction

The present application is directed to a switching network for receiving and transmitting data packets having both frames and messages which utilizes a ring for messages and an associated crossbar switch for frames.

Background of the Invention

In a switching network, all receiving channels (or ports) route data to a switching fabric, which sends switches the data, which is normally in the form of data packets either of uniform or variable length, to a specific destination port. A data packet may include both frames which consist of relatively long strings of data bytes for example 40 to 64 bytes and larger, and also include messages which consist of small entities of, for example 4, 8, or 12 bytes. Such small entity messages might include formats of broadcast flow control, back pressure/feed forward messages, linked table configuration, write or read formats and other similar formats. Input ports are connected to output ports by a well known crossbar connection matrix. Such crossbar matrices typically reside on a die where there may be 64 ports and each port has a data bus of 16 signal lines. Thus, with a total of 2,048 signal lines, the crossbar switches are silicon resource intensive. In other words, to efficiently utilize this silicon resource (that is the silicon die on which the crossbar switch is integrated), it is very inefficient to transmit small entity messages (that is 4, 8, or 12 bytes, for example, as discussed above) through the crossbar switch. It is more efficient, rather, to transmit frame size packet portions which range from 40 to 64 bytes and greater.

Ring networks have also been suggested for data transfer. See IEEE 802.5 standard. However, this is used in a computer network where a computer must first catch a token and then attach a "message" to it.

Object and Summary of the Invention

It is a general object of the present invention to provide switching network using a message ring to receive and transmit messages.

In accordance with the above object, there is provided a switching network for receiving and transmitting data packets having both frames which consist of relatively long strings of bytes of, for example, 40 to 64 bytes, and messages which consist of small entities of, for example 4, 8, 12 bytes comprising a ring of data ports. Crossbar means connect the ports for switching the frames from an input port to an output port. Ring means successively interconnect one port to an adjacent port in the ring for forming the ring for passing the messages from an input port, successively through intermediate ports to a destination output port. The frames and messages are processed simultaneously.

Brief Description of Drawings

Fig. 1 is a simplified block diagram of switching apparatus embodying the present invention.

Fig. 2 is a circuit schematic of message ring architecture embodying the present invention.

Fig. 3 is a block diagram illustrating the operation of Fig. 2.

Fig. 4 is a circuit schematic illustrating the operation of Fig. 2.

Fig. 5 is a flow chart illustrating the operation of Fig. 2.

Fig. 1 is an overall diagram of a switching apparatus which include as an essential component the switching network of the present invention. Specifically, there are 8 switch elements designated SE0 through SE7. Each of these switch elements have 64 input and 64 output lines. There are equivalent numbers of switching networks in each of the switching elements. The overall switching apparatus in Fig. 1 is also disclosed in a co-pending application entitled "Switching Apparatus For High Speed Channels Using Multiple Parallel Lower Speed Channels While Maintaining Data Rates" Filed October 3, 2001, Serial No. 09/971,097 (Attorney Docket No. 6979/13).

Referring in general to the operation of the switching apparatus of Fig. 1, there are a number of ingress source ports 11 numbered 0 through 63 each receiving data packets from, for example, a framer which normally puts together such packet, at a rate of 10 Gbps. The ingress

ports 11 include a TM (traffic manager) and a communications processor (CP) and are labeled CP/TM. Each ingress source port has an 8-line output port, each individually coupled to an input port of switch elements SE0 through SE7 which together create a so-called switching fabric. In turn, the eight switch elements each with 64 input ports and 64 output ports are similarly connected on an output side to egress ports 12 also designated CP/TM which have 8-line inputs and are numbered 0 through 63. The combination of the 64 ingress ports and 64 egress ports make up a 64 port full duplex port.

Again, as on the input side, each output port of a switch element has a direct serial link to one of the CP/TMs or egress port units. Then the egress ports 12 are coupled into, for example, a high speed channel network (e.g., fiber optic) to transmit data at a 10 Gbps rate in a manner similar to the incoming data, but with the data having been rerouted to a selected destination port. Finally, as indicated in Fig. 1, the high input and output data rates of 10 Gbps cannot normally be sustained separately by the switch elements SE0 through SE7 which as indicated are limited to a lower data rate of 2.5 Gbps.

Fig. 2 illustrates a combined crossbar switch 510 with a message ring 550 having a number of input ports nominally designated 500a through 500h. From a practical standpoint, in the context of the present invention, there is one input port (and one output port) for each of the 64 lines shown in, for example, switching element SE0 in Fig. 1. Thus, the circuit of Fig. 2 is an integrated portion of each of the switching elements SE0 and SE7 of Fig. 1. Each port may either be a source, that is input, or destination port depending on the nature and the location of the switching element. The switching network of Fig. 2 forms a typical crossbar switch (as discussed above) where the internal crossbar switch unit 510 receives from the various input ports 500a through 500h, data streams from the various communications processors/traffic managers 0 through 63 illustrated in Fig. 1. Referring briefly to Fig. 3, each port of the switching network of Fig. 2, is associated with a parser/FIFO illustrated in dash outline 20 in Fig. 2 and shown in greater detail in Fig. 3. On line 21, data packets are routed to or from a CP/TM at a 2.5 Gbps rate. A parser 22 identifies whether the portion of the data packet is a message or frame and then respectively sends it to a frame FIFO 23 or a message FIFO 24 (FIFO being an abbreviation for First In First Out memory). Then, on the input/output lines 26, 27 of the respective FIFOs, the frame or message data is input to a port or node 500a-500h (one of the 64 ports) and processed or switched as determined by the ring controller 520 and the clock 560.

If a frame is being routed to a desired destination port, the crossbar switch 510 operates in a normal manner where, for example, data would be input into the node 500h directly switched to the crossbar switch 510 and then immediately switched to the desired destination port. As discussed above, to perform this switching with a small entity message would be both inefficient and unduly congest the crossbar switch. Thus, if a message that is in place or queued up in message FIFO 24 as illustrated in Fig. 3, it is inserted a particular node or port (assuming the port has no other data present in it at the moment) and then passed successively through intermediate ports via the interconnecting lines 600 between ports until the final destination port is reached. Thus, the interconnecting lines 600 between the ports 500a-500h form the message ring 550. Under the control of clock 560, messages are moved from one available port to the next for every clock pulse.

In order to avoid conflict with the crossbar switch, however, each port 500a-500H includes, as illustrated in Fig. 4, a gate 31 (nominally of the AND type) which buffers a data input 32 to an output register 33 which is connected to, for example, a port 500h under the control of line 34 from the controller 520. This prevents conflict with the simultaneous crossbar switching of the same switching network as illustrated in Fig. 2.

Fig. 5 is a flow chart illustrating the operation of Figs. 3, 4 and 5. In the step 200 a data packet is analyzed by the parser 22 and it is determined whether it is a message or frame. Then in step 210, if it's a frame, it is routed in the conventional manner through the crossbar switch as discussed above. If a message is placed in a message in queue in step 220 (as also illustrated in Fig. 3) it is handled in a first in, first out (FIFO) manner. In step 230 the message is inserted into one of the ports or nodes of the message ring, that is 500a-500h, and is also given a message ring destination identifier in step 240. It is passed from port to port in step 250 under the control of the clock 560 and the gate unit of Fig. 4. In step 260 the question is asked if the message is at its destination port. If no, it is passed to the next port in step 270 but if yes as indicated in Fig. 5, it is placed in a message out queue in step 280. And as illustrated in Fig. 3, the message out queue is a message FIFO which is operating in an output manner.

Thus, messages do not pass through the crossbar 510 as illustrated in Fig. 3 but instead they are passed directly through the message ring from port to port. Thus, congestion of the crossbar switch is minimized.

In summary, a switching network for receiving and transmitting data packets having both frames and messages is provided by the use of a message ring.